Ferrofluid-Mediated Lattice Processor Architecture for Permutable Application Optimization Post-Manufacture and Enhanced Heat Dissipation

21 December 2024
Simon Edwards
Research Acceleration Initiative

Introduction

Although application-specific processor architectures have been in use for quite some time, chips which seek to integrate a variety of ASICs on a single chip (ideal for performance) do not allow for any great flexibility in what percentage of the chip surface is optimized for conventional, machine learning or high-floating-point functionality.

Abstract

I propose a fundamentally new approach which can allow for application-specific optimization in variable proportion on a single chip wherein the percentage of the surface area of the chip optimized for one function or another can be altered after the initial manufacture of the chip.

Rather than pipelines being placed directly in-line with transistors (on the same plane,) I propose that a ferrofluid layer of nanoscopic thickness be emplaced between pairs of transistor layers. Each pair of computing layers would have, therefore, a shared pipeline configuration wherein the same pipelines are shared by each pair of layers. Interference between the two layers may be prevented by directing current in opposing directions in each of the individual layers in a pair and by permitting flow through transistors in only a single direction.

The ferrofluid would; with the aid of an external magnetic field; form pipelines which could be varied in order to change the relative configuration of the pipelines. For high-floating-point applications, parallel, linear pipelines would be desired, but a branching structure would be desired for more conventional computing. Greater inter-connectivity would be emphasized for Machine Learning applications.

The ferrofluid layers, which would be manifold and would be sandwiched between each pair of transistor layers, would serve not only to form the basis of conductive pipelines connecting the transistors, but would serve the double-purpose enhancing thermal dissipation. As both pipelines and transistors are heat-generating, keeping transistor layers and pipeline layers separate except at needed junctions would serve to prevent overheating.

The transistors would necessarily be configured as a lattice so as to enable maximal flexibility as defined by the relative structure of the pipelines, which would be fundamentally fluid-based.

Conclusion

Entities with variable computing needs would likely find such processors useful. With a sufficient level of control over pipeline architecture, extremely specific, repetitive tasks might be hardware-optimized which do not constitute a sufficient share of the overall processing workload to justify the manufacture of discrete chips, especially when the tasks in question can change on a daily basis. Although contractors might be asked to build application-specific chips for these specific applications, as the application for which the chips are to be used might be inferred from knowledge of the architecture of the chips, the ability to alter the system of inter-connectivity of transistor architectures in-house would greatly enhance operational security.